

*Permissible Suspension and Evidence Resistance*

This chapter surveys recent accounts of the epistemic permissibility of suspended judgement in an attempt to thereby identify the normative resources required for explaining the epistemically problematic nature of evidence resistance. Since paradigmatic cases of evidence resistance involve belief suspension on propositions that are well supported by evidence, such as vaccine safety and climate change, the literature on permissible suspension seems to be a straightforward starting point for my investigation: after all, any plausible view of permissible suspension will have to predict epistemic impermissibility in these paradigmatic resistance cases. I look at three extant accounts of permissible suspension – a simple knowledge-based account, a virtue-based account, and a respect-based account – and argue that they fail to provide the needed resources for this project. Further on, the chapter identifies the source of the said difficulties and gestures towards a better way forward.

**4.1 Proper Suspension: The Simple Knowledge-Based View**

It is widely agreed that justification – be it moral, prudential, epistemic, etc. – is defeasible. For instance, suppose that you justifiably head towards High Street on a Sunday because you wish for a new pair of shoes, but as you're walking, I tell you that you left your wallet at home. In this case, you have a defeater for your (prudential) justification for going into town. Should you continue on your way, your action will no longer be (prudentially) justified. Similarly, suppose that you (epistemically) justifiably believe that the structure you are looking at is a barn. Suppose, further, that I tell you that most of the things that look like barns are actually fakes. In this case, you have a defeater for your belief that the structure you are looking at is a barn. If you continue to hold this belief, your belief is no longer justified.

While it is widely agreed that epistemic justification is defeasible, and much ink in epistemology has been spilled on the issue of the defeasibility of the justification of positive doxastic attitudes, such as beliefs and credences, very little has been said about the justification of suspension and about its defeasibility conditions. However, paradigmatic cases of evidence resistance involve epistemically unjustified belief suspension on propositions that are well supported by evidence. Of course, this need not be the case: one can be evidence resistant by simply not taking up evidence and not updating properly in the light thereof, without thereby being suspended on the issue: indeed, it may even be that one hosts a fairly high credence that  $p$  is the case while, at the same time, resisting genuine evidence for  $p$  due to, for instance, bias against a particular source.

However, given that paradigmatic cases of evidence resistance involve defeated suspension, work on permissible suspension seems like an excellent place to start an inquiry into resistance impermissibility. As such, the fact that the amount of work done to date on the justification and defeasibility of suspension is relatively minimal poses a problem.

I take the most notable views of permissible suspension on the market to be the following: (1) the simple knowledge-based account (e.g. often implied but not often explicitly defended in Williamson 2000, Sutton 2005, 2007, Hawthorne and Srinivasan 2013), (2) the virtue-based account (Sosa 2021), and (3) the respect-based account (Miracchi 2017, Sylvan and Lord 2021, 2022).

The knowledge-based account, while not being explicitly defended in many places, follows straightforwardly from endorsing a biconditional knowledge norm of belief (defended most notably in Williamson (2000), but also in, e.g., Sutton (2005, 2007), Hawthorne and Srinivasan (2013), and Littlejohn (2020)), according to which belief is epistemically permissible just in case it is knowledge. If non-knowledgeable belief is impermissible, it follows that one should suspend belief in cases in which one does not know.

To put my cards on the table from the very beginning: I take a simple knowledge-based account of permissible suspension to suffer from insurmountable in-principle difficulties in dealing with resistance cases. On a view like this, one is permissibly suspended on  $p$  just in case one does not know that  $p$ . The main trouble for any account along these lines comes from the sufficiency claim: since knowledge implies belief, the sufficiency of lack of knowledge for permissible suspension claim implies that suspension implies permissible suspension. After all, should one not know in

virtue of refusing to believe, on a view that takes lack of knowledge to be enough for permissible suspension, one would thereby be diagnosed as permissibly suspended. 'x implies permissible x' is not a great result for suspension or, more generally, for any x, subject to any sorts of normativity. Furthermore, and even more interestingly, upon closer scrutiny, the account has further problematic results: since belief is predicted to be impermissible in cases of lack of knowledge due to lack of belief, suspension of belief will not only be merely permissible, but even obligatory. To see this, take a paradigmatic case of evidence resistance. I falsely believe that not-*p*: it is not the case that climate change is happening. All experts in climate change come and tell me that *p*. In virtue of politically motivated bias, I refuse to believe them and suspend on the issue. At this stage, the simple knowledge account's diagnosis is that indeed I should suspend, since I don't know that climate change is happening – even though the reason why I don't know is because I refuse to believe it. As such, on the simple knowledge-based view, suspension implies obligatory suspension.

I take these considerations to signal insurmountable difficulties for a simple knowledge-based view of permissible suspension.

Williamson is sensitive to this problem when he writes:

A Pyrrhonist sceptic may hope to comply vacuously with all three norms [(N) Believe only what you know, (DN) Be the sort of person who is disposed to believe only what one knows, (ODN) Do the thing that a person who is disposed to believe only what she knows would do] by having a general disposition never to believe anything. If one has no beliefs, then a fortiori one has no untrue beliefs, no beliefs that fail to constitute knowledge, no beliefs that are improbable on one's evidence, no inconsistent beliefs, and so on. The Pyrrhonist, if such a person is possible, complies with all three norms even in the sceptical scenario. [...] Non-sceptics may find little to admire in the Pyrrhonist's self-imposed ignorance, especially when that ignorance concerns the needs of others. There may be positive norms for knowledge, such as a norm enjoining knowledge-gathering in various circumstances, and so positive as well as negative norms for beliefs. (Williamson forthcominga)

I agree with Williamson that what is needed here is input from positive epistemology. This book attempts to do just this, and it does so in keeping with the knowledge-first picture – albeit not with a simple version thereof. Chapter 9 develops a view according to which proper suspension has to do with what one is in a position to know.

Before this, however, in what follows, I will look at a prominent recent defence of a virtue-theoretic account of permissible suspension from Ernie Sosa.

#### 4.2 Sosa on Virtue, Telic Normativity, and Suspension

Very little has been said in the literature about the justification of suspension and about its defeasibility conditions. Ernie Sosa's most recent book (2021) supplies this lack: Sosa offers a comprehensive virtue-theoretic account of the nature and normativity of suspension in terms of the nature and telic normativity of agential attempts more generally.

In what follows, I first briefly outline the position and take issue with some details of its normative structure. In particular, I argue that Sosa's telic normativity is in need of normative expansion if it is to accommodate the defeasibility of justification to suspend. Further on, I consider several paths for developing Sosa's view to accommodate this datum and argue that we can find the needed resources in general telic normativity.

Sosa's virtue epistemology is a normative framework for the evaluation of attempts (henceforth also 'telic normativity'). Attempts have constitutive aims. As a result, we can ask whether or not a given attempt is successful. We can also ask whether a given attempt is competent (i.e. produced by an ability to attain the attempt's aim). Finally, we can ask whether a given attempt is apt (i.e. successful because competent).

Virtue epistemologists standardly take beliefs to be attempts that have truth and/or knowledge as their constitutive aims. Given that this is so, we can ask whether beliefs are successful (i.e. whether they are true). In addition, we can also ask whether they are competent (i.e. whether they are produced by an ability to believe truly) and whether they are apt (i.e. true because competent).

According to Sosa, the above gives us the basic account for first-order evaluations of attempts. Crucially, however, Sosa does not take this to be the whole story. Rather, he countenances two further types of aptness alongside first-order aptness, or 'animal' aptness as Sosa calls it. These additional types of aptness are 'reflective' and 'full' aptness. Attaining these further types of aptness requires accurate and indeed apt attempt at a higher order, in addition to animal aptness. In a nutshell, the thought is that attempts will rise to these higher levels of aptness only if, alongside animal aptness, one has aptly ascertained that one's attempt is free from any relevant risk that one may be running: one must have arrived at an apt awareness that one's attempt would be apt. While animal aptness in conjunction with apt risk assessment will be enough for reflective aptness, full aptness additionally requires that first- and second-order aptness are connected in the right way: one must be guided to animal aptness by one's reflectively apt risk assessment.

It comes to light that there are a number of normative properties that attempts can enjoy. Crucially, according to Sosa, full aptness enjoys special status among these properties. More specifically, according to Sosa, full aptness is the fully desirable status for attempts, and attempts fall short unless they attain full aptness. Moreover, he is also clear that this claim holds with full generality: any attempt attains fully desirable status qua attempt if and only if it is fully apt; and it falls short qua attempt if and only if it isn't.

According to Sosa, various psychological categories – most importantly, guessing, belief, and judgement – are species of affirmation and, as a result, attempts. (Sosa's main interest is with affirmations with a specifically epistemic aim that, at a minimum, involves truth.) While Sosa countenances a variety of psychological categories with epistemic aims, his main focus is on judgement (and judgemental belief). Judgement differs from other psychological categories in that it has a particularly robust epistemic aim: judgement aims not only at truth, but at aptness. To understand this normative requirement on judgement, Sosa asks us to consider Diana, the huntress: as Diana surveys a landscape in search of game, she may see prey in the distance (in good light and calm wind). If a shot is too risky, it is ill-advised. A shot, then, can attain quality in it being well rather than negligently selected. An aiming, then, is assessable by reference to how likely it is to succeed (relative to one's possession of the pertinent competence), so as to avoid recklessness, and it is also assessable by reference to how negligent (or not) it may be.

Similarly, according to Sosa, for a judgement to be apt, more is required than merely apt affirmation. What is needed for apt judgement is that one is guided to aptness by apt risk assessment. An apt judgement is a fully apt affirmation.

Where does suspension fit in this picture? After all, telic normativity is a normativity of attempts, but isn't suspension a paradigm of something that is not an attempt, but rather an instance of forbearing from attempting?

To answer this question, Sosa introduces a distinction between two varieties of intentional forbearing:

*Narrow-scope:* (Forbearing from X'ing) in the endeavour to attain a given aim A.

*Broad-scope:* Forbearing from (X'ing in the endeavour to attain a given aim A).

According to Sosa, the first, narrow-scope variety of forbearing pertains to telic normativity proper: the forbearing is done with the domain-

internal aim in view. The second, in contrast, is domain-external forbearing, in that the agent who forbears in this sense does not attempt to reach the central aim of the domain in question to begin with: whether to engage in a domain is not a question within the domain itself. In that, broad-scope forbearing, according to Sosa, does not make the proper subject of telic normativity.

To see the place of forbearing in the normativity of attempts, consider Diana again. Diana's archery shots can be more or less well selected. When she spots some prey, Diana can properly aim as follows: to make an attempt on that target if and only if the attempt would succeed aptly. Accordingly, there are two ways in which Diana can fall short with regard to this aim: she could make an attempt on the target when she would not succeed aptly – because, maybe, the shot would be too risky, given the wind. But she could also fail in her attempt by failing to make an attempt (on the target) when one would succeed aptly.

So, in a nutshell, according to Sosa, narrow-scope forbearing is itself an attempt with an aim: that of attempting if and only if the attempt would succeed aptly. This is the place of forbearing in telic normativity.

How does this translate to epistemology? Again, just like with normativity in general, Sosa thinks that it is only narrow-scope forbearing that is of internal interest to the theory of knowledge proper in that it is aimed at the epistemic goal of attaining aptness. More specifically, Sosa thinks that epistemic narrow-scope forbearing is what constitutes deliberative suspension of judgement, which is an attempt in its own right, one that shares with judging an epistemically distinctive aim: the aim of affirming alethically (positively or negatively) if and only if that affirming would be apt (and otherwise suspend). Conversely, on Sosa's view, one properly suspends belief on a question if and only if one suspends based sufficiently on one's lack of the competence required in order to answer that question aptly (2021, 85).

In contrast, broad-scope forbearing, according to Sosa, is the stuff of intellectual ethics (i.e. it pertains to the question as to whether to engage in inquiry as to whether  $p$  to begin with). In this sense, it is external to the theory of knowledge proper. Here is Sosa:

Whether to engage in a certain domain is not generally a question *within* that domain. Telic assessment within a domain assesses mainly the pursuit of aims proper to that domain. An exhausted tennis competitor *may* of course properly consider whether to default, but this is not a decision assessable *within* the sport. When you sense a heart attack in progress and quit for that reason, this is not a decision assessable by athletic criteria in the

domain of tennis. Whether to keep on playing is not a tennis decision; it is a life decision. (2021, 66, emphases in original)

Similarly, Sosa thinks that epistemic broad-scope forbearing is tantamount to non-deliberative suspension of judgement. It is also an intentional forbearing from alethic affirmation (both positive and negative), but it is not aimed at apt judgement as to whether  $p$ ; it derives rather from omitting inquiry into the question as to whether  $p$  to begin with, whether the refusal is implicit or consciously explicit. As such, norms governing broad-scope forbearing will be norms of intellectual ethics, not epistemic norms proper:

Broad-scope forbearing [i.e. not taking up a question] is not a standing *within* the domain of inquiry into a particular question, wherein it would be subject to the epistemic assessment of attempts that are potentially knowledge-constitutive. (2021, 70–71, emphasis in original)

To bring my first worry into clear view, I'd like to start with a particular case of evidence resistance involving ignored defeat. To take a variation on a famous example, consider the case of a scientist, Gabriel, who doesn't believe anything his female colleagues say, because he is a sexist (Lackey 2018). Now suppose Gabriel carries out two experiments to test his hypothesis that  $p$ . Experiment §1 strongly supports that  $p$ . Experiment §2 strongly supports that not- $p$ . The scientist comes to suspend on  $p$  on this basis. Suppose, next, that a female colleague of his, Dana, discovers a serious flaw with experiment §2, which she points out to Gabriel. Due to sexist bias, Gabriel discounts Dana's word and maintains his suspension on  $p$ . This is a paradigm case of higher-order defeat: after Dana's testimony that  $q$ : 'there is a flaw in experiment §2', Gabriel's suspension on  $p$  is no longer justified.

What does Sosa's account have to say about this case? It would seem that, for all we have been told so far, telic normativity does not have the resources to accommodate the result that Gabriel is not justified to suspend. Rather, Gabriel's failure will, at best, be categorised as pertaining to intellectual ethics. To see this, note that Gabriel never takes up the question as to whether  $q$  to begin with due to his sexist bias. As such, since no attempt at apt judgement is made, the suspension at stake in the case of  $q$  will have to be classified as non-deliberative suspension. If that is so, however, its normative properties will not have the capacity to affect the normative properties of Gabriel's suspension on  $p$  either: after all, even if present, normative failure outside the domain of theory of knowledge proper need not affect domain-internal normative properties. Even if

Gabriel's suspension on  $q$  is impermissible on non-epistemic grounds, it cannot affect the epistemic permissibility of Gabriel's suspension on  $p$ .

Recall also that, on Sosa's view, suspension is permissible insofar as it is sufficiently based on one's lack of the competence required in order to answer that question aptly. It is easy to see that this account predicts, against intuition, that it is permissible for Gabriel to suspend based on his sexism-generated lack of competence to believe aptly what Dana tells him.

Now, it is worth mentioning that there may be an easy way for Sosa out of this case: one thing he could do is insist that Gabriel does, in fact – albeit implicitly – inquire into whether  $q$  by simply hearing the testimony from Dana. After all, Sosa's notion of inquiry is a very 'light' one, whereby the mere monitoring of one's environment counts as such. If so, Gabriel will count as having epistemically impermissibly suspended on  $q$ , since, in the course of his (implicit) inquiry into whether  $q$ , he missed the opportunity to affirm aptly that  $q$ .

That said, the route back to problems for Sosa's account is quite short from here. To see this, note that we can easily tweak the case following a recipe that will be familiar by now, such that Gabriel doesn't even hear that Dana told him that  $q$ . Once more, suppose that Gabriel simply zones out whenever a female colleague talks to him. As a result, Gabriel didn't even register that Dana told him that there is a problem with his experiment. In this case, Gabriel's epistemic behaviour is no better than in the original case. If anything, it's worse. Most importantly for present purposes, the case is equally one of testimonial injustice and one of defeat. Once Gabriel is told about the flaw in his experiment, Gabriel's suspension on  $p$  is no longer justified. The fact that Gabriel didn't bother to listen does not improve his situation vis-à-vis the original case on either count.

Sosa has not discussed the issue of normative defeat directly. However, in *Epistemic Explanations* (2021), he has started theorising about negligence within his virtue-epistemological framework. Most importantly for present purposes, he suggests that negligence may preclude competent performance. In particular, negligent failure to inquire may preclude competent judgement. If so, we could maybe avail ourselves of this normative resource to explain how negligent failure to inquire may preclude competent suspension as well.

Note, first, that cases of normative defeat do plausibly count as cases of negligent failure to inquire. Consider the case of Gabriel, the sexist scientist, once more. Gabriel is told by his female colleague, Dana, that there is a flaw in one of the experiments that led him to suspend on  $p$ , but



Gabriel doesn't even listen. Isn't this a prime example of a negligent failure to engage with the question as to what he was told? If Sosa is right and negligent failure to inquire precludes competent judgement, then presumably it also precludes competent suspension. Given that justified suspension is competent suspension, we get the desired result that Gabriel is not justified in his suspension.

Unfortunately, there remains a fly in the ointment: negligence is itself a normative property. If your failure to inquire into whether  $p$  is negligent, then you didn't inquire into whether  $p$ , although you *should* have. Crucially, while one may agree that we need to understand normative defeat in terms of violations of the norms requiring us inquire, the task Sosa faces is to offer an account of these norms within the scope of theory of knowledge proper – rather than intellectual ethics. For virtue epistemologists like Sosa, this means offering an account that is available to virtue epistemology. Since the kind of negligence that precludes justified suspension is a normative epistemic property, this means that what we need is a substantive account of the kind of negligence that precludes justified suspension in terms of the abilities or other resources available in the theoretical machinery of Sosa's framework. To say that cases of external defeat are cases in which competent suspension is precluded by negligent failures to inquire gives us a way of identifying the task that we are facing, but not yet a way of accomplishing it.

Unfortunately, there is an in-principle reason to worry that it will not be trivial to accomplish this task, given Sosa's framework. To see why, note again that Sosa conceives of telic normativity as the normativity of attempts: whether an attempt is successful, competent, or apt presupposes that an attempt was made. In this way, telic normativity presupposes that the agent has made an attempt. As a result, whether or not the agent *should* make an attempt is not assessable in terms of Sosa's telic normativity of attempts. Recall also that, to make sense of norms requiring us to inquire, Sosa distinguishes between the epistemic normativity of the theory of knowledge (i.e. telic normativity) and the broader normativity of inquiry. Obligations to inquire fall into the broader normativity of inquiry, which pertains to intellectual ethics.

The trouble is that Sosa's suggestion that negligence may preclude competent judgement is hard to square with the claims above. To see this, let's return to the case of the sexist scientist once more. Recall that the thought was that when Gabriel doesn't even listen to Dana, he falls foul of negligent failure to engage with the question of what Dana tells him. But negligence is normative: to be negligent is to fail to do certain things that one should have done. In particular, the way in which Gabriel is negligent

here is that he fails to take up the question of what Dana tells him, even though he should have done so.

We are now in a position to see the in-principle problem for Sosa. If Gabriel's negligence consists in his failure to take up the question whether  $q$ , even though he should have done so, his failure does not fall within the normativity proper to the theory of knowledge, but into the broader epistemic normativity of inquiry. As a result, it is now hard to see how his negligence may preclude deliberative competent suspension on  $p$ . After all, deliberative competent suspension does fall within the normativity proper to the theory of knowledge. At the same time, this normativity is autonomous and protected from incursion of extraneous normativity, including that of the broader normativity of intellectual ethics. It looks as though accounting for cases of normative defeat in terms of negligence that we are envisaging is not available to Sosa after all, at least not provided that the rest of his theory stays put.

Sosa (2021) does offer an account of the kind of negligence that is at stake in the cases discussed. He considers a case in which you are adding numbers via mental arithmetic. If the set of numbers you are adding is sufficiently large, you will not be sufficiently reliable to arrive at a competent belief about the sum. Suppose you are still sufficiently reliable but barely so. At the same time, you have a calculator ready at hand, which would keep you safely above the relevant reliability threshold. If you insist on mental arithmetic here, Sosa argues, you fall foul of negligence.

With the case in play, let's move on to Sosa's view of negligence. Here is the crucial passage:

I am suggesting that negligence is a failure of competence, that one proceeds inappropriately in performing as one does if one *should* have taken the steps by not taking which one is negligent. One is then to blame (in the negligence mode) for not having taken those steps. [...] *Competent attainment of aptness requires availing yourself of sufficiently available means that would enable a more reliable assessment of your first order aptness and competence.* If there are no such means, then there is no such negligence, and no such incompetence. In such a circumstance, the agent might then be able to determine with sufficient competence that they are in a position to proceed competently enough on the first order. (2021, 64, emphases in original)

Sosa's key idea is that if you can assess your first-order competence by more reliable means but fail to do so, then you are negligent. In particular, you fall foul of a kind of negligence that precludes what he calls the competent attainment of aptness.

Most importantly for present purposes, given that competent suspension requires that one suspends based sufficiently on one's lack of

the competence required in order to answer the question aptly, negligence precludes competent suspension: sexist scientist Gabriel does have sufficiently available means that would enable a more reliable assessment of the aptness of his suspension – Dana’s testimony. Since he ignores it, Gabriel will count as a negligent suspender.

The problem with this account of negligence, however, is that it is too strong: it makes negligence, and hence defeat, too easy to come by. To see this, consider a case in which I ask my flatmate, who is currently in the kitchen, whether we have any milk left. He tells me that we do. Now, I do have several more reliable means of assessing my first-order competence available to me. For instance, I could go to the kitchen and have a look myself. Crucially, however, failure to avail myself of these means doesn’t make me negligent. And, most importantly for present purposes, it doesn’t preclude my judgement that there is milk in the fridge from being competent.

Sosa’s account of negligence is insufficiently normative. What matters, according to Sosa, is the *availability of alternative means* that would lead to a more reliable assessment of first-order aptness and competence. However, the difference maker is normative, not descriptive: what matters is not (only) whether one has alternative means available that would have led one to a more reliable assessment of first-order aptness and competence, but (also) whether one *should* have availed oneself of these means. In the case of the sexist scientist, he should have taken the woman’s testimony into account in assessing the credentials of hypothesis *p*. Similarly, Mary should not have ignored all of the evidence suggesting that her husband is having an affair. In contrast, in the milk case, it is not the case that I should have had a look myself.

Sosa’s account of the normativity of negligence in terms of available alternative means doesn’t work unless we add that the available means are means one should have availed oneself of. Crucially, it is precisely this ‘should’ that we wanted to explain in virtue-epistemological terms. We are thus back to square one once more.

In more recent work (Sosa 2022), and in reply to my worries, Sosa goes into more detail about the normativity of negligence. On this novel view, availing oneself of more reliable alternative means is only normative in cases in which the initial performance was lacking to begin with. Looking left and right, for instance, is something you could easily enough have done and should have done before crossing the street. Omitting to do so is a case of negligence. This refurbished view straightforwardly takes care of the flatmate case, since that is a case of permissible belief formation.

The question that arises is: how does this account of negligence fit with Sosa's general normative picture?

According to Sosa, such lacks and omissions can diminish the agency of an agent and can thus be negatively assessable telically even though they are not attempts. Since they are attributable to the agent, they thereby speak to their skill. When you cross that London street without looking, the failure to look is presumably attributable to you – it speaks badly of your competence.

Here is, however, the worry I have for this reply: Sosa's telic epistemic normative picture is a reliabilist normative picture. As such, Sosa's competences admit for failure even when they are manifest – indeed, any reliabilism will predict this. Consider, now, one-off cases of suspension on  $p$  in spite of easily available evidence for  $p$ : say that Gabriel, the scientist, is not a sexist, and indeed he's a fantastically competent epistemic agent, but he simply fails this time around to give the weight it deserves to his colleague's testimony. This failure will not speak against his competence – indeed, it cannot: failures are predicted in a reliabilist picture about competence. If this failure does not speak to Gabriel's competence, however, Sosa's picture is left without resources to predict it as epistemically relevantly negligent – since, on his view, epistemically relevant negligence is competence-relevant negligence. As such, Sosa's picture will lack the needed resources to accommodate cases of one-off impermissible suspension by competence-manifesting epistemic agents.

One way around this problem<sup>1</sup> would be to derive the impermissibility of negligence straight from the success condition involved in Sosa's picture – in the case of judgement, from knowledge. I agree: this is exactly the route I will take in Chapter 9.

### 4.3 The Respect-Based View of Permissible Suspension

More recently, Lisa Miracchi (2017) and Kurt Sylvan and Errol Lord (2022) have proposed more nuanced virtue-theoretic views of permissible suspension. The views are more nuanced than their straightforward knowledge-based and virtue-based competition in the following ways: in contrast to the simple knowledge-based account, these views do not make

<sup>1</sup> Suggested by Ernie Sosa in conversation. Another route Ernie seemed partial to in conversation is to take some varieties of failure to imply lack of competence manifestation. The difficulty with this approach will be finding a principled way to distinguish between varieties of failure that are compatible with competence manifestation and varieties of failure that are not.

direct appeal to the absence of epistemic value to explain the permissibility of suspension. In contrast also to Sosa's classic virtue-theoretic approach, their accounts conceive of the permissibility of suspension in terms of manifesting *respect* for the epistemic value in question – be it truth or knowledge – rather than in terms of value-conduciveness.

Henceforth, I will run with Miracchi's view, due to hers being the first such proposal on the market. According to Miracchi, epistemic virtue not only involves aiming to get what is fundamentally valuable (knowledge, on her knowledge-first virtue epistemological picture), but also involves respecting the aim of getting what is fundamentally valuable. Respecting the aim of getting knowledge, she claims, is derivative from the fundamental aim. On Miracchi's view, suspension is permissible just in case it manifests respect for the aim of knowing. In this, the features that make epistemic rational assessment applicable to suspension are derivative from the features that make such assessment applicable to beliefs. Here is Miracchi:

When an agent generally can be characterized as aiming to A (e.g. aiming to know), we can understand withholding from or omitting a performance of A-ing as manifesting a kind of practical respect for what it takes to A. The agent manifests this respect precisely by not endeavoring to A. (2017, 433)

This account is well equipped to deal nicely with many paradigmatic cases of evidence resistance: it seems right that, in not uptalking relevant, easily available evidence, our resistant agents in the cases under discussion fail to manifest respect for what it takes to know, which, indeed, casts doubt on their knowing competences. This suggests that the sufficiency direction of Miracchi's view holds.

I think Miracchi's derivative, respect-based account makes significant progress in understanding the paradigmatic virtue-theoretic flaws involved in impermissible suspension. Unfortunately, I don't think it will quite get us there as an *analysis* of permissible suspension due to considerations having to do with the strength of the respect condition. Here is why: one worry one might have for Miracchi's view follows the necessity direction – isn't a requirement of manifesting respect for knowledge going to be too strong for permissible suspension? After all, one might think, paradigmatically bad believers – conspiracy theorists, wishful thinkers, bullshitting politicians – can also, on occasion, suspend properly on everyday matters – for instance, on whether there is milk in the fridge. Think back also to the Pyrrhonist sceptic: this agent, one might think, is the paradigmatic example of a bad suspender, with little respect for knowledge.

Compatibly, should the Pyrrhonist (or conspiracy theorist, etc.) suspend on whether there's milk in the fridge based on, for example, conflicting testimony, their suspension would be justified.

Miracchi is aware that the respect condition needs to be fairly weak to do the job. She writes:

In the case of knowledge, [respect] can be in the form of competently taking yourself not to have (or be able to have) sufficient evidence to settle  $p$ ?, and so intentionally adopting a settled attitude of suspension on  $p$ ?. But it can also be just a matter of having, on the first order, a competence to know. Having good dispositions to withhold is essential to possessing competences to know. (Otherwise, except in very special environments, the reliability condition would fail.) [. . .] When a person withholds in these ways, we can say that she demonstrates her competence to know without exercising it. (2017, 433)

As such, at least on the face of it, making the respect condition very weak may afford Miracchi the following way to accommodate one-off permissible suspension cases: paradigmatically bad believers can also, on occasion, manifest respect for knowledge simply in suspending in ways that demonstrate their having a competence to know – even if, in the vast majority of their epistemic walks of life, they fail to do so, and thus fail to either manifest or demonstrate their said competence. Systematic failure to manifest a competence need not imply one does not have the said competence. I can have a competence to play the piano fabulously and still fail to do so whenever I get to it because, say, I am distracted. Insofar as the competence exists, however, it allows for one-off manifestations of competence, as well as for one-off demonstrations of competence. I can still play the piano fabulously on the few occasions when I care to do so (and thereby manifest my competence); I can also abstain from playing when I ascertain, for instance, that the conditions are not favourable (too dark), and thus I demonstrate my competence in forbearing from playing. Similarly, one could argue, bad believers and the Pyrrhonist may still have a competence to know and thus both manifest and demonstrate it on occasion. So far, so good.

Unfortunately, upon closer inspection, this response is not available to Miracchi. Miracchi's view is a reliabilist one. This will present her account with a strength dilemma related to cases of bad believers and what is involved in having a competence to know: according to Miracchi, the having of the said competence has to do with it being sufficiently objectively likely that whenever the sub-personal cognitive mechanisms that constitute the basis of the competence are operative, the conditions constitutive of knowledge obtain. It is easy to see that the case of bad believers

presents an account like this with a strength dilemma: make the sufficiency threshold for the objective likelihood at stake too high and bad believers don't count as having the competence; if so, the account cannot explain one-off cases in which these people suspend permissibly. Conversely, make the likelihood threshold too low and the view is not a plausible reliabilism anymore.

The second reason why a reliabilist view like Miracchi's will struggle with permissible suspension goes, to some extent, back to the discussion of Sosa's negligence condition in the following sense: at the core of any reliabilist account lies an allowance for fallibility. My competences to know can fail to generate knowledge even when exercised; these are what Miracchi (2015) calls degenerate exercises of competence. When this happens, on Miracchi's view, one forms justified beliefs. The question that arises now is: if competences to know can be manifest in belief-formation episodes that fail to result in knowledge, it seems plausible that we should take the corresponding competence-demonstrating respect for knowledge to be manifest even in instances in which there is failure of permissible suspension. After all, since the competence is fallible, it should also be possible for it to fail in this way: I am an excellent epistemic agent, I evaluate my evidence thoroughly, and I come to suspend on the issue; unfortunately, this is one of the (otherwise very few) instances in which I am wrong about what my evidence supports. My suspension, intuitively, demonstrates my (fallible) competence to know; nevertheless, my suspension is unjustified. Note that this is also plausible on general grounds having to do with how manifesting respect works more generally: it doesn't always lead to success in treating the respected party in the right way due to, for example, moral bad luck. But if manifesting respect for knowledge allows for instances that lead to improper suspension, we are left without resources to explain what exactly is going wrong in these cases.<sup>2</sup>

#### 4.4 Gesturing towards a Better Way

I share with several of the people discussed so far a commitment to a telic normative structure. In what follows, I thus want to go back to Ernie

<sup>2</sup> Finally, a similar problem arises (even more straightforwardly) for the Sylvan and Lord (2022) Kantian incarnation of the view. According to them, what it is for A's doxastic reaction (be it belief, suspension, etc.) D to respect the truth is for D to be the output of an exercise of a competence to react for normative epistemic reasons. Since the only plausible way to conceive of the competence at stake in creatures like us is to allow for its fallibility, the view will have trouble explaining the impermissibility of suspension in one-off cases of failure to properly assess one's normative reasons.

Sosa's view, for illustrative purposes, and gesture at a different way to accommodate the normative defeasibility of suspension within general telic normativity. In particular, I will suggest that what is needed is to enlarge the normative remit of epistemic telic normativity in line with plausible normative facts about general telic normativity.

Recall that we have identified two in-principle problems with the virtue-theoretic account of epistemically permissible suspension under discussion. First, Sosa's epistemic telic normativity is the normativity of attempts, but in the cases under discussion no attempt is being made to begin with: the defeating evidence is totally ignored. As such, what we need is an account that accommodates attempts that *should* have been made.

Second, on Sosa's account of suspension, one properly suspends belief on a question if one suspends based sufficiently on one's lack of the competence required in order to judge aptly. However, many of the protagonists in the cases we have looked at do lack the relevant competences: the sexist scientist, for instance, is not a competent uptaker of testimony from women due to his sexist bias. He does suspend based on his (sexism-induced) lack of competence to judge aptly. What seems to matter, then, is not whether one misses a competence or not, but rather whether one *should* have had the competence to begin with.

Both of these points suggest that we need more normative resources than epistemic telic normativity, as put forth by Sosa, provides: for a correct account of justified suspension, we need to be able to also assess (at least some) attempts that *should* be made and competences one *should* have within theory of knowledge proper, rather than merely at the level of intellectual ethics.

At the same time, of course, some 'shoulds' governing attempts and competences will fall outside of theory of knowledge proper indeed, falling squarely within the remit of intellectual ethics. The question as to whether I should or should not know more about mathematics, the geography of oceans, and the workings of the human lungs than I presently do will not concern the theory of knowledge, and the corresponding normative failures – should I exhibit them – will not defeat my justification for my current beliefs and suspensions.<sup>3</sup>

If all of this is right, it would seem that what needs to be done is that we must move the border between the theory of knowledge proper and intellectual ethics, such that we allow some 'shoulds' governing attempts

<sup>3</sup> But see Chapters 8 and 9 for more discussion of epistemic obligations to inquire.



and competences to fall on the side of theory of knowledge while others remain squarely within intellectual ethics.

I will begin by discussing shoulds governing attempts. First, to see why it is independently plausible that attempts that should have been made can be domain-internal, let's go back to Diana, the huntress: Diana's archery shots can be more or less well selected. We have seen that Sosa agrees that there are two ways in which Diana can fall short with regard to her aim to succeed aptly: she could make an attempt on the target when she would not succeed aptly – because, maybe, the shot would be too risky, given the wind. But she could also fail in her attempt by failing to make an attempt (on the target) when one would succeed aptly. There are thus two types of meta-competence failure that Diana can display: failure to assess risk properly, but also failure to assess opportunity properly. When going back to epistemology, the latter failure is the stuff of unjustified suspension: a failure to judge (affirmatively or negatively) when one would have judged aptly (independently of whether one attempted to do so or not).

Now, here is one question: why think that Diana's failure of the second kind (opportunity assessment failure) is conditional upon her making any attempts (including attempting to shoot and including attempts to shoot if and only if the shot is apt) to begin with? Why think that this should pertain to the normativity of extant attempts rather than to the normativity of attempts that should have been made? After all, it is plausibly constitutive of the huntress's professional role that she should make hunting attempts, including attempting to shoot if and only if the shot is apt. A huntress that fails to make any hunting attempts is a rubbish huntress. The meta-competence to assess risks and opportunities in Diana's case is not attempt-conditional. It is also, at the same time, not domain-external: the question is not whether Diana should become a huntress to begin with – that's, of course, the stuff of professional ethics. Rather, what is going on is that, in her capacity as a huntress, Diana shoulders shoulds pertaining to attempts she should make, not just shoulds governing the ones she does make. Indeed, plausibly, these shoulds are constitutive of what it is to be a huntress to begin with.

If this is so, on pain of losing the analogy, we should expect that the normativity internal to the domain of the theory of knowledge proper follows suit: there will be attempts that the epistemic agent should make, given that the opportunity arises to judge or suspend aptly as a result of making said attempts. Epistemic agents who will ignore easy opportunities by not even attempting will be rubbish epistemic agents, just like

huntresses who don't bother to take easy targets, or who don't even bother to assess shooting opportunities, are rubbish huntresses.

There is, of course, an important disanalogy between the two cases: one can choose not to be a huntress. It's harder for agents like us, with our cognitive capacities, to choose not to be epistemic agents. If so, the domain-external question – should I engage in epistemic endeavours? – does not even arise for us: we just can't help it. What room is there left, then, on this picture, for questions of intellectual ethics?

Note that Diana is not an ideal huntress: there are limits to the amount of opportunities she can take. Should she find herself in a forest filled with thousands of easily available targets, she can only reasonably be expected to make a limited number of attempts. Likely, she will be normatively constrained to shoot at the most readily available targets. For the rest, it's up to her: she can't attempt to shoot at all of them, so it's up to other normative considerations not pertaining to the domain of hunting to decide which to go for. Maybe Diana has moral concerns against shooting cubs; maybe she has prudential interests in favour of shooting valuable prey; in all of these cases, these domain-external normative considerations will guide her choices.

Our epistemic environment is a bit like the forest filled with too many shooting opportunities. We have plenty of opportunities to judge aptly about thousands of things just as we walk down the street. We can't take them all: we are psychologically limited creatures. Some we should (epistemically) take: I should form the belief that there's a building before me when it's in plain sight; I should believe the testimony of others, absent defeat; and so on.

For the rest, there will be many opportunities that I just can't take because of the limited kind of being that I aim: there's a limited number of attempts at apt judgement I can make. That's why whether I decide to study maths is a question of intellectual ethics, guided by prudential, moral, and other non-epistemic normative constraints. Epistemology only asks that I take the easiest of opportunities that lie right in front of me, just like hunting only asks that Diana makes attempts at the easy targets.

This concludes my discussion of attempts one (epistemically) should make: they correspond to (easy) epistemic opportunities one should take because one would thereby aptly judge. The should at stake is internal to the epistemic domain because it pertains to what it is to be a good epistemic agent to begin with.

How about cases in which you lack the relevant competence to begin with, although you should have had it? Recall that the case of the sexist

scientist is plausibly like that: he does suspend based on his lack of competence to judge aptly, which, in turn, is triggered by his sexism: he can't give the woman the credibility she deserves. Is this failure also going to be epistemic domain-internal? After all, by stipulation, the sexist scientist does not miss an opportunity to believe aptly, since he lacks the competence to properly assess the woman's credibility to begin with.

I suggest that we step away from epistemology once more, go back to cases of general telic normativity, and ask the question: is it plausible to think that there are norms internal to the domain of hunting that regulate what competences huntresses should have? I think the answer is clearly 'yes'. Indeed, it is arguable that these are norms that are constitutive of the domain: huntresses should, at a minimum, be able to spot her prey, shoot, and hit it with some degree of reliability in normal environmental conditions. Huntresses who lack these basic abilities are rubbish huntresses; indeed, if they lack them all, they may no longer count as huntresses to begin with. And this is not the stuff of professional ethics but rather is constitutively normative of the domain of hunting itself.

Similarly, I want to suggest that epistemic agents who lack competences that are constitutive of the kind of epistemic agents that we are rubbish epistemic agents by the light of normativity internal to the epistemic domain itself. Sexists, hallucinators, and wishful thinkers alike are in breach of epistemic norms proper. This explains why the normativity of competences one should (epistemically) have can affect the normative status of one's epistemic attempts.

This section did not have the ambition to develop a full account of permissible suspension within a more inclusive theory that takes competences to matter for justification, but rather to gesture towards the resources that are needed if we are to make progress on this front. Chapter 9 will develop an account along these lines of permissible suspension that places centre stage the notion of being in a position to know.

#### 4.5 Conclusion

This chapter looked at what I take to be the three most prominent accounts of permissible suspension recently put forth in the literature: the simple knowledge-based account, the simple virtue-based account, and the respect-based account. I have argued that the simple knowledge-based view faces insurmountable difficulties, in that it implies a highly problematic claim: that suspension implies obligatory suspension. I have then argued that Sosa's epistemic telic normativity is in need of normative

expansion or else it cannot deal with normative defeat. I have also looked at a respect-centric variety of virtue theory about permissible suspension from Miracchi and Sylvan and Lord, and I argued that, while the knowledge/truth respect condition on permissible suspension that they favour may deal exceptionally well with paradigmatic cases of impermissible suspension, it is both too strong and too weak to be useful for an analysis thereof. Finally, on a positive note, I have suggested that, in line with general telic normativity, we should conceive of epistemic telic normativity as also concerning attempts we should have made, as well as competences we should have had.